



Objective versus Relative Risk in Privacy Decision Making: A Replication Study from Germany

Sebastian Hermes

Technical University of Munich
sebastian.hermes@tum.de

Luis Hillebrand

Technical University of Munich
luis.hillebrand@tum.de

Markus Böhm

Technical University of Munich
markus.boehm@tum.de

Jan Bauer

Technical University of Munich
jan.bauer@tum.de

Helmut Krcmar

Technical University of Munich
helmut.krcmar@tum.de

Abstract:

This study reinvestigates the effects of normative and behavioral factors on privacy decision making by conducting a methodological replication of Adjerid, Peer, and Acquisti (2018). While the normative perspective regards consumers with stable preferences making rational choices, the behavioral perspective regards consumers with unstable preferences making irrational choices due to heuristics and biases. In three experiments, we demonstrate that normative and behavioral factors influence hypothetical but not actual choice. Our results, therefore, confirm the findings of the original study that objective differences in privacy protections influence hypothetical choice. However, in contrast to the original study, we found that relative changes in privacy protection did not influence actual but hypothetical disclosure as well. We argue that individuals have developed a stronger disposition toward privacy since the original study and that our German student sample represents a more privacy-sensitive case than the American Amazon Mechanical Turk sample. As a consequence, participants may have not been willing to indicate their true choice in the actual setting. In other words, effects may exist in the actual setting, but may not be elicitable from privacy-sensitive individuals. Future research is encouraged to explore other biases and the moderating effect of disposition to privacy.

Keywords: Privacy, privacy decision making, behavioral economics, prospect theory, methodological replication

The manuscript was received 7/15/2020 and was with the authors 8 months for 2 revisions.

1 Introduction

The desire to understand consumer data and privacy preferences has sparked interest in research, practice, and legislation in equal measure. While firms need personal information to personalize their services and improve the effectiveness of their marketing campaigns (Farahat & Bailey, 2012), policymakers seek to reduce consumer harm and protect social and economic welfare from privacy violations. Understanding the factors and mechanisms of consumer privacy decision making has therefore become a vital topic across multiple research domains. However, previous IS research has focused on either normative (rational decision making) or behavioral (irrational decision making) aspects to account for changes in privacy choices but has neglected to explore both perspectives simultaneously.

One work aiming to understand how behavioral and normative aspects simultaneously influence privacy decision making is the study of Adjerid et al. (2018) published in *Management Information Systems Quarterly*: “Beyond the Privacy Paradox: Objective versus Relative Risk in Privacy Decision Making”. The study incorporates a behavioral perspective of privacy decision making by building upon prospect theory (Kahneman & Tversky, 1979) in which heuristics and biases are accounted for. The authors operationalize their objective by investigating how differing degrees of privacy protection influence consumers’ willingness to disclose personal information. In three experiments, the authors compare the impact of objective risk of disclosure and relative perceptions of risk of disclosure on both hypothetical and actual information disclosure in English-speaking subjects recruited on Amazon Mechanical Turk and Prolific. The three experiments conducted in the original study were driven by normative and behavioral theories such as the privacy calculus (Dinev & Hart, 2006) for the former and prospect theory (Kahneman & Tversky, 1979) for the latter. While normative factors refer to rational and stable preferences of utility-maximizing agents (Mullainathan & Thaler, 2000), behavioral factors refer to unstable and irrational preferences that stem from limitations in consumers’ cognitive ability such as reference dependencies and heuristics in the case of a survey’s look and feel (John, Acquisti, & Loewenstein, 2011).

Since previous IS privacy research struggles to simultaneously study the impact of normative and behavioral factors and, to the best of our knowledge, no work that replicates the study of Adjerid et al. (2018) exists, this paper aims to fill this gap. Therefore, we conduct a methodological replication wherein the theories, methods, and hypotheses are adopted from the original study of Adjerid et al. (2018). There are two reasons why we selected this paper for replication. First, it focuses on behavioral factors (reference dependency), which remains a scarce endeavor in the IS community, although some initial work exists and the subfield continues to develop (Herrmann, Kundisch, & Rahman, 2014; Keith, Babb, & Lowry, 2014). Second, this paper adopts an experimental methodology, which is beneficial for replication, because experiments allow for a greater degree of control than other behavioral approaches (Dennis & Valacich, 2015). We now present the research overview and hypotheses adopted from the original study (see Figure 1 and Table 1).

Objective Differences in Privacy Protection

Consumer privacy decision making can be affected by changes in perceived privacy benefits and risks. For example, individuals might provide personal information if they expect to receive more personalized products or services (Adjerid et al., 2018; Ansari & Mela, 2003). Similarly, individuals might conceal information if they believe their disclosure will pose significant risks (Dinev & Hart, 2006; Malhotra, Kim, & Agarwal, 2004) such as price discrimination (Viswanathan, Kuruzovich, Gosain, & Agarwal, 2007). Following this line of thought, Adjerid et al. (2018) propose that privacy protections influence privacy decision making via their impact on perceived risks of information misuse. Hence, *Hypothesis 1* proposes that manipulating normative factors such as objective levels of privacy protection will affect privacy decision making such as information disclosure (Table 1).

Relative Changes in Privacy Protection

Previous work on behavioral factors indicates that privacy decision making can also be relative in nature (Acquisti, John, & Loewenstein, 2012). For example, heuristics, biases, and emotions such as joy and fear have been found to influence how consumers perceive privacy protection and privacy risk (H. Li, Sarathy, & Xu, 2011). A fruitful theoretical lens for analyzing the relative nature of privacy decision making has been offered by Kahneman and Tversky (1979). The authors introduced Prospect Theory in 1979 and challenged

the expected utility theory developed by Neumann and Morgenstern (1944) by demonstrating that individuals also make irrational choices, such as making decisions based on perceived gains instead of perceived losses. However, the proposition that individuals' decision making can also be influenced by reference points is of particular interest for this study. Outcomes above or below the reference point are considered as gains or losses. Acquisti, John, and Loewenstein (2013), for example, demonstrate that individuals are more likely to keep their data private if their data has already been kept private compared to individuals whose data has not been kept private in the first place. Hence, *Hypothesis 2* proposes that behavioral factors such as relative changes in privacy protection influence privacy decision making such as information disclosure (Table 1).

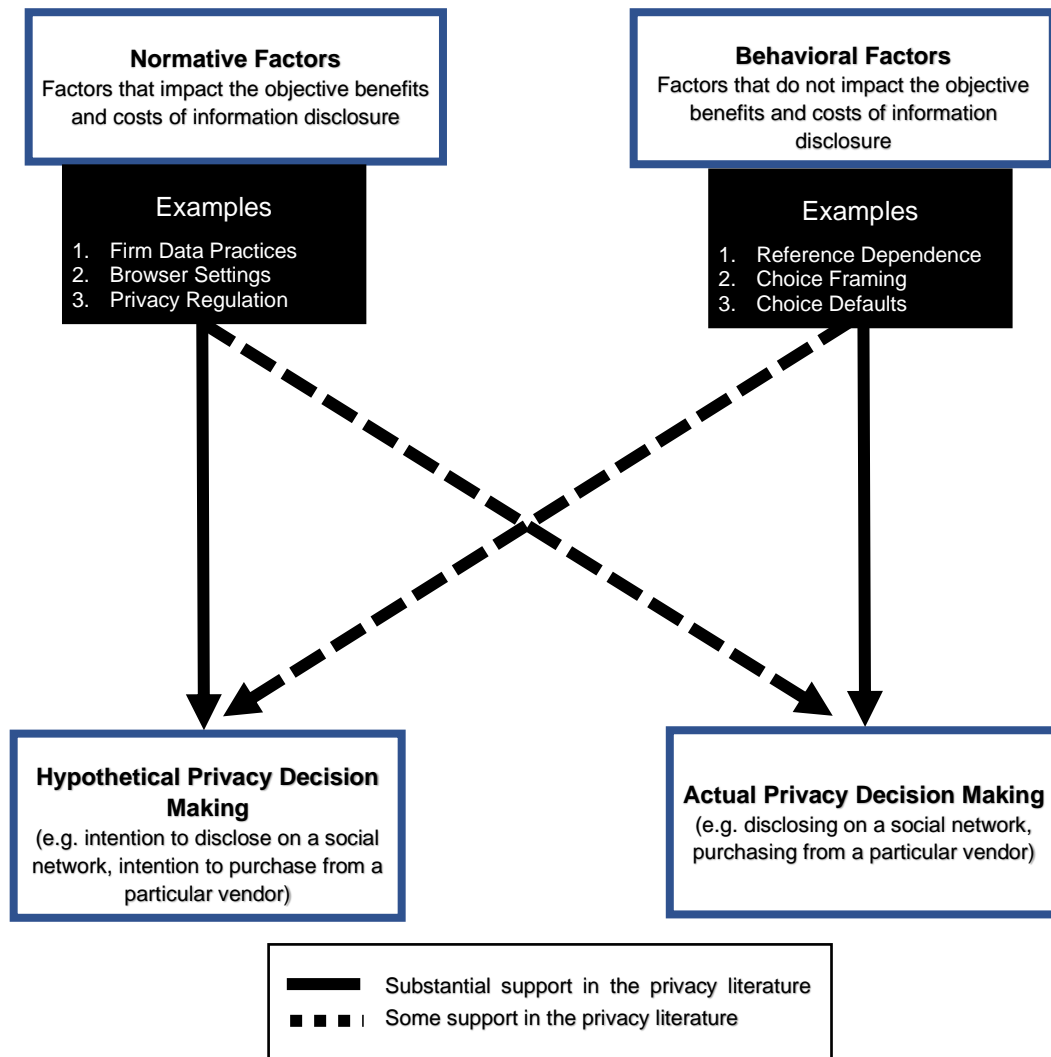


Figure 1. Research Overview (Adjerid et al., 2018)

Privacy Decision Making in Actual Versus Hypothetical Disclosure Contexts

Although comprehensive evidence exists for the normative and behavioral perspectives, it remains unclear how normative and behavioral factors influence hypothetical and actual information disclosure (Adjerid et al., 2018). On the one side, there may be no difference between the two and on the other, the influence of both factors may vary across hypothetical and actual disclosure settings. If normative factors vary across both disclosure settings, this would constitute a hypothetical bias, indicating a gap between behavioral

intentions and actual behavior. LaPiere (1934) was the first to observe this bias by studying race prejudice. The author found that 92% of the respondents stated that they would not accommodate members of the Chinese race, while in reality, 95% actually did accommodate them. Hence, hypothetical bias refers to the phenomenon that individuals may indicate an intention that they fail to live up to in practice. Empirical studies support that this phenomenon is prevalent (Ajzen, Brown, & Carvajal, 2004; FeldmanHall et al., 2012). Adjerid et al. (2018) further argue that the intention-behavior gap occurs due to more positive attitudes toward a behavior in a hypothetical rather than actual disclosure setting. In other words, if positive attitudes toward protecting privacy exist, these attitudes are going to influence hypothetical rather than actual disclosure. Hence, *Hypothesis 3* proposes that normative factors are stronger in hypothetical compared to actual disclosure settings.

The influence of behavioral factors may also vary across hypothetical and actual disclosure settings. Previous work suggests that behavioral factors have at least some impact in actual disclosure settings (Knetsch, Tang, & Thaler, 2001) and may play a stronger role in actual than in hypothetical settings. Kang and Camerer (2013) and Loewenstein (2000), for example, show that individuals are state-dependent and fail to anticipate the actual choices they will make in future hot states (state in which they are impacted by visceral drivers such as hunger) when considering the same choice context hypothetically. Put simply, individuals in a hot state do not fully understand how much their behavior is influenced by their current state and individuals in cold states find it difficult to imagine themselves in hot states. Translated to privacy decision making, these results indicate that individuals may be unable to anticipate their hot state (e.g., how privacy choice contexts are framed) when considering hypothetical disclosures as opposed to actual disclosures. Hence, *Hypothesis 4* proposes that behavioral factors are weaker in hypothetical than in actual disclosure settings.

Table 1. Research Hypotheses	
H1	Changes in objective levels of privacy protection will affect disclosure: lower levels of privacy protection will lead to lower levels of disclosure of personal information.
H2	One's relative perception of the level of privacy protection will influence individual privacy decision making: levels of privacy protection perceived to be higher relative to a reference point will result in higher levels of disclosure of personal information.
H3	The impact of normative factors (i.e., objective changes in privacy protection) will be stronger on hypothetical intentions to disclose compared to actual disclosures.
H4	The impact of behavioral factors (i.e., relative changes in privacy protection) will be weaker on hypothetical intentions to disclose compared to actual disclosures.

However, in contrast to the original study, we do not draw upon a sample of American and English-speaking participants recruited on Amazon Mechanical Turk and Prolific. Instead, we recruited German students and ask them to forward the survey to their families, friends, and colleagues from work. This sample provides the opportunity to identify whether the results presented in the original study are generalizable to populations beyond those in the USA and English-speaking realms and whether the results hold multiple years later.

We chose to focus on German students and their social entourage for three reasons. First, we expect that Americans and German perceive privacy differently (Fromholz, 2000) and that Germans' high levels of uncertainty avoidance (Hofstede-Insights, 2019; Hofstede, 2001) translates into high levels of need for privacy which can affect information disclosure (Y. Li, 2014). Second, the public and scholarly debate of the General Data Protection Regulation (GDPR) in Europe is likely to have increased Germans' privacy sensitivity (especially among students). Moreover, we expect that the increasing exposure to privacy scandals (Clement, 2019a, 2019b) has also increased the privacy sensitivity of our sample and that such individuals will be more restrictive about information disclosure compared to the individuals of the original study. Third, we focus on students and their social entourage to address the limitations that come with pure student samples and thereby aimed for more robust and generalizable results.

The remainder of this paper is organized as follows: First, we present our methodology and summarize and discuss our results; next, we outline practical and theoretical implications; and finally, we highlight limitations for our work and illustrate fruitful avenues for future research.

2 Methodology

This replication study follows the methodology and the three experiments conducted in the original study of Adjerid et al. (2018). In the original study, online pools from Amazon Mechanical Turk and Prolific Academic were used to gain a sufficient sample size. Unlike the original study, this study conducts the three experiments with German students in the field of business administration. The students voluntarily participated and received a 10-point bonus in their course for doing so. To increase the sample population, we requested the students to ask family members, friends, and colleagues from work to participate in the study. Students obtained 1 point for each referral who completed the survey (no more than 3 points were granted in total). We instructed the students to not provide any further information about the experiments to their friends, colleagues, or family before they shared the survey to ensure students did not influence response behavior. We also instructed students to not recruit participants from the course. Participants had 2 weeks (June 14, 2019, to June 28, 2019) to complete the experiments. The students were randomly assigned to one experiment based on their last name¹. To match the sample size of the original study, we matched more students to Experiment 3, since the sample size of this experiment in the original study was larger than that in the other two experiments. LimeSurvey, an online statistical survey web app, was used to create the experiments.

Our variables are exact replications of the original study. The only exceptions are the questions' intrusiveness (a control variable to assess the effect of questions that had been judged in Acquisti et al. (2012) as highly intrusive on disclosure) and the survey's visual design (a control variable to assess the effect of the survey's visual design on disclosure). We did not consider those two control variables since they were only used in Experiment 2 of the original study. The original study revealed that the survey's visual design has no effect on disclosure. Intrusiveness, however, had mainly a negative effect confirming prior working on information sensitivity (Malhotra et al., 2004). The manipulations that we are interested in are captured by the different groups of participants, which differ regarding their privacy protection levels. We evaluated the impact of manipulations on non-repeating dependent variables (e.g., privacy concerns and protection satisfaction) to assess whether the manipulations led to different perceptions of privacy protection. To this end, we used t-tests and chi-square tests. For all experiments, we relied on either actual or hypothetical willingness to disclose as dependent variables. Both disclosure settings asked participants to make a series of disclosure decisions. To appropriately analyze this experimental setup, we conducted random-effects regression analysis. We considered a participant-specific random effect.

3 Experiment 1: Hypothetical information disclosure

3.1 Methodology

For the first experiment, we randomly assigned participants to each treatment. The experiment investigated hypothetical willingness to disclose personal information. Participants were told at the beginning of the study that they had to complete two separate surveys (named Survey A and B), which included hypothetical sensitive and ethical questions. Then, in the first part of Experiment 1 participants received either high or low levels of privacy protection. After protection recall questions and manipulation checks (see Appendix B, Table B1), both groups answered ten questions about their hypothetical willingness to disclose ethically sensitive information (see Appendix B, Table B2.). We added an attention check between Survey A and Survey B.

¹ The randomization based on the last name was successful. There were no significant differences in the demographic distributions, in Experiment 1, for age ($t(232.31) = .229, p = .81$), and gender ($X^2(3, N = 235) = .488, p = .485$), in Experiment 2, for age ($F(3,415) = .145, p = .93$) and gender ($X^2(3, N = 419) = 2.412, p = .491$), and in Experiment 3, for age ($F(1,684) = .92, p = .338$) and gender ($X^2(3, N = 686) = 1.679, p = .641$).

In the second part of Experiment 1, both groups perceived either an increase or decrease in relative privacy protection, while the actual privacy level was held constant between the two groups in the second part (at a medium privacy level). Protection recall questions and manipulation checks were presented again. Both groups had to answer the same questions about their willingness to share sensitive information as in the first part. Finally, there were follow-up questions about general online privacy concerns and demographic questions (age and gender). Figure 2 illustrates the process of Experiment 1 in a flow chart.

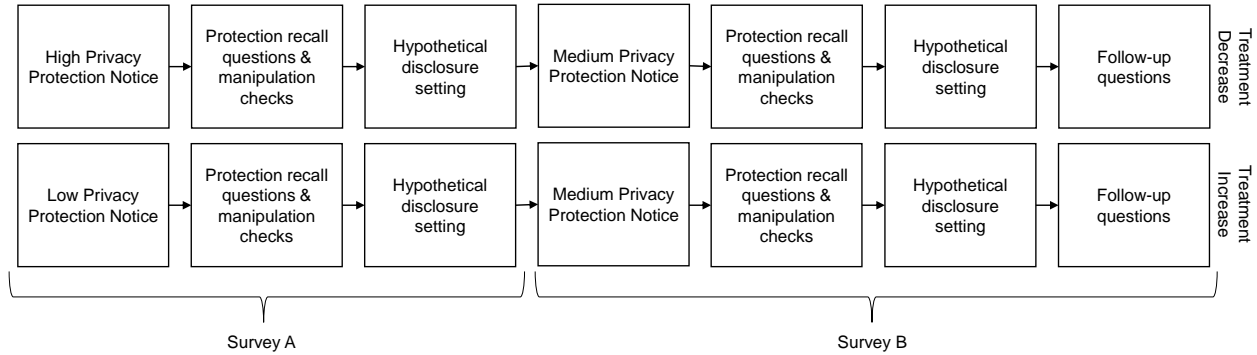


Figure 2: Flow Chart of Experiment 1

We included some reverse answer options for the Likert scales of the sensitive questions to improve the validation of the study by comparing two additional groups with and without reverse answer options. We compared the two groups with a t-test to assure that both indicated similar responses. To indicate the privacy protection level of each survey, we used a graphical representation (see Appendix B, Figure B1) as described in the original study.

Table A1 in Appendix A shows the demographic data of all participants who successfully completed Experiment 1. In the original study, the total sample size for the first experiment was 221 (37.56% female and a mean age of 29.16, SD of age is 9.76) (Adjerid et al., 2018). Our sample size was 235 (46.6 % female and a mean age of 24.96, SD of age is 6.96).

3.2 Results

By and large, we found that participants were able to understand the privacy protection notices provided in the experiment. Although for Surveys A and B only 67.2% and 40.0% correctly recalled at least four of the five dimensions, our manipulation of objective risk was indeed effective in influencing the perception of privacy protection levels in the first survey (Survey A). Participants in the high protection group were significantly more satisfied with those protections ($M_{High} = 3.84$, $M_{Low} = 2.85$), $t(219.65) = 7.12$, $p < .001$, $d = 1.64$, significantly less concerned about privacy ($M_{High} = 2.65$, $M_{Low} = 3.58$), $t(232.89) = -5.7284$, $p < .001$, $d = -1.64$, and significantly less concerned about harm that would come to them as a result of disclosing personal information ($M_{High} = 2.50$, $M_{Low} = 3.00$), $t(232.35) = -3.11$, $p < .01$, $d = -1$ (see Table 2).

Table 2. Experiment 1, Summary Results						
Conditions	Survey A			Survey B		
	High Protection	Low Protection	p-value	Increasing	Decreasing	p-value
Our results						
Privacy Concern	2.65	3.58	p<.001	3.05	3.42	p<.016
Protection Satisfaction	3.84	2.85	p<.001	3.12	2.90	p<.15
Harm Perception	2.50	3.00	p<.001	2.82	3.08	p<.068
Original results						
Privacy Concern	2.39	3.87	p<.001	2.76	3.29	p<.01
Protection Satisfaction	3.36	1.56	p<.001	2.86	2.41	p<.01
Harm Perception	2.86	4.02	p<.001	3.37	3.68	p = .04

We used random-effects regression to estimate the effects of the manipulation. Participants reported their likelihood of disclosure for a given question on a five-item scale (1 = "Very Unlikely" to disclose, 5 = "Very Likely" to disclose). We found that the objective differences in privacy protection levels in Survey A had a significant effect on participants' predicted behavior. Participants that were given a low level of privacy protection said that they were significantly less likely ($\beta_{\text{Low}} = -.32$, $p < .01$) to disclose personal information (Table 3, column 1). This was consistent ($\beta_{\text{Low}} = -.32$, $p < .001$) when question type (descriptive versus ethical), participants' age, and gender were included as control variables (Table 3, column 2). These results provide strong support for the hypothesis that objective risk will affect consumer privacy choices in a hypothetical disclosure setting (H1 is supported).

For the second survey (Survey B), which had an objectively identical medium level of privacy protection for both conditions, participants in the increasing-protection condition reported being more, but not significantly more, satisfied with the protections provided ($M_{\text{Inc}} = 3.12$, $M_{\text{Dec}} = 2.90$), $t(228.44) = 1.42$, $p = .16$, $d = 0.40$, not significantly less concerned that their responses might be used in ways that could harm them ($M_{\text{Inc}} = 2.82$, $M_{\text{Dec}} = 3.08$), $t(232.55) = -1.83$, $p = .07$, $d = -0.47$, but significantly less concerned about privacy ($M_{\text{Inc}} = 3.05$, $M_{\text{Dec}} = 3.42$), $t(232.92) = -2.42$, $p < .05$, $d = -0.47$. Different from the original study, the relative change in privacy protection in Survey B *did* have a significant effect on participants' predicted disclosure behavior. Specifically, we found that increasing privacy protection *did* have a significant effect ($\beta_{\text{Increasing}} = .28$, $p < .05$) on overall predicted disclosure levels (Table 3, column 3). This result is robust ($\beta_{\text{Increasing}} = .28$, $p < .05$) when controls for question type and participant age and gender were included (Table 3, column 4). Hence, in the hypothetical disclosure setting, our results support the hypothesis that the relative perception of privacy protection influences disclosure behavior (H2 is supported).

Variables	Admission (1 Very Unlikely – 5 Very Likely)							
	Our results				Original Study			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
Low Protection	-0.324**	-0.320**			-0.669**	-0.650**		
	(0.118)	(0.118)			(0.120)	(0.118)		
Increasing			0.278*	0.282*			0.0925	0.109
			(0.129)	(0.128)			(0.123)	(0.120)
Descriptive		0.053		0.096*		-0.494**		-0.565**
		(0.049)		(0.047)		(0.0607)		(0.0601)
Age		-0.012		-0.016		-0.0132*		-0.0100
		(0.008)		(0.009)		(0.00651)		(0.00680)
Gender		-0.126		-0.129		0.130		0.196
		(0.118)		(0.129)		(0.124)		(0.129)
Constant	3.229**	3.666**	2.776**	3.402	3.631**	4.173**	3.328**	3.772**
	(0.084)	(0.249)	(0.092)	(0.268)	(0.0701)	(0.229)	(0.0784)	(0.249)
Observations	2,350	2,350	2,350	2,350	2,210	2,210	2,210	2,210
Number of id	235	235	235	235	221	221	221	221

Robust standard errors in parentheses; **p < 0.01, *p < 0.05, +p < 0.1.

3.3 Discussion

Our results suggest that both objective differences and relative changes in privacy protection levels influence privacy perception. More precisely, we found that perceived risk of harm, satisfaction with privacy measures, and privacy concerns were significantly different between objectively high and low privacy protection levels. We also found evidence that privacy concerns are significantly different when privacy protections increase or decrease. However, perceived risk of harm and privacy satisfaction were not significantly different in relative privacy protection changes.

We found effects on hypothetical information disclosure for both objective and relative changes in privacy protection levels. Thus, supporting both H1 and H2. Experiment 1 also provides initial support for H3, as the impact of normative factors (objective change) on hypothetical intentions to disclose information may be more pronounced in hypothetical settings. However, given that there was no comparison with data on actual disclosure, this is only suggestive. Experiment 1 does not seem to support H4, since we identified that relative changes have a significant effect on hypothetical disclosure. The subsequent experiment investigated how normative and behavioral factors influence actual disclosure.

4 Experiment 2: Actual information disclosure

4.1 Methodology

The second experiment was conducted with a different group of students than the first experiment. Unlike the first experiment, where hypothetical disclosure was examined, this experiment focused on actual disclosures while manipulating objective and relative changes in privacy protection. The survey was a 2 x 2 between-subject design and participants were randomly assigned to one of the four groups. Participants were manipulated in such a way that they perceived either an increase, decrease, or the same level of privacy protection for two different surveys (named Survey A and B). At the beginning of the experiment,

participants were told that they would have to participate in two separate surveys and would receive confirmation codes for each survey via email. The confirmation code was needed to prove that they collected data so that participants could receive the course bonus.

At the beginning of the first survey, participants provided their email address, age, and gender. Thereafter, the privacy protection notice was displayed, which conveyed either a high or low level of privacy protection (see Appendix C, Table C1). As in the original study, we used the same text-based privacy level notices. We included additional protection recall questions and manipulation checks (see Appendix C, Table C2). In the next step participants had to answer six questions about ethically questionable behavior (Acquisti et al., 2012) (see Appendix C, Table C3.). As in the first experiment, we placed an attention check between the first and the second survey. Then, we included a reverse answer scale for one of the six personal questions.

Identical to the original study, the second part of Experiment 2 looked and felt different from the first part (see Appendix C, Figure C1 and Figure C2). Again, all participants had to provide their email address and some demographic information (age and gender). After the privacy protection notice and protection recall question and manipulation checks, participants were asked six different questions about ethically questionable activities (Acquisti et al., 2012) (see Appendix C, Table C3.). At the end, some exit questions were presented (e.g., whether the privacy level had changed between the two parts) (see Appendix C, Table C4)². Figure 3 illustrates the process of Experiment 2 in a flow chart.

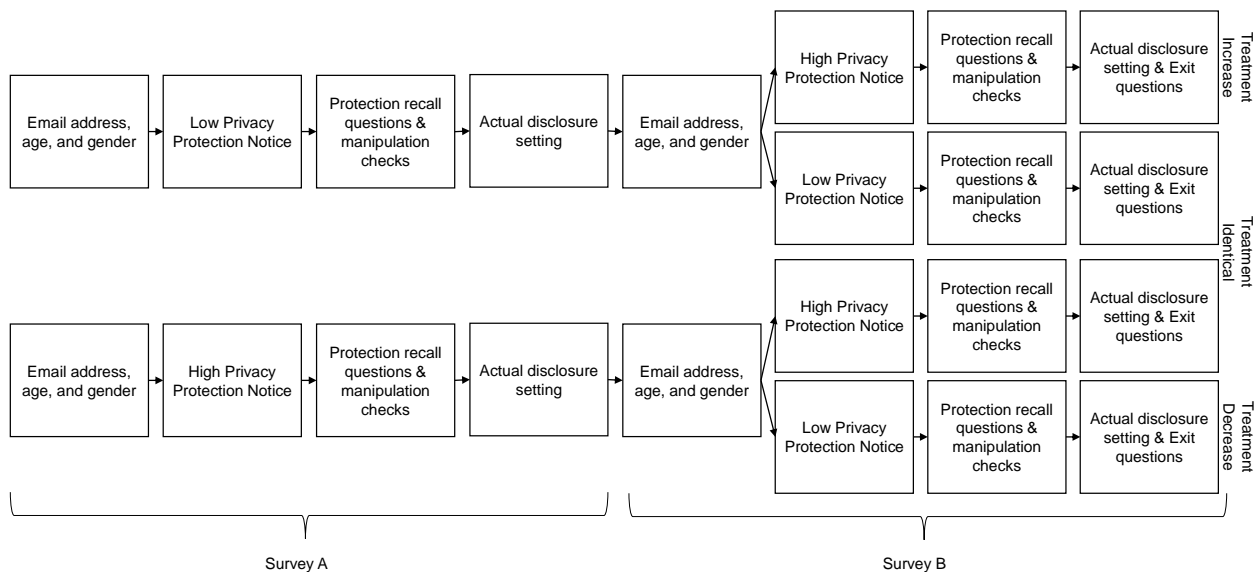


Figure 3: Flow Chart of Experiment 2

Table A2 in Appendix A shows the demographic data of all participants who successfully completed the experiment. In the original study, the total sample size for the second experiment was 415 (51.61% female and a mean age of 31.27, SD of age is 10.72) (Adjerid et al., 2018). Our sample size was 412 (50.02% female and a mean age of 25.12, SD of age is 9.02).

4.2 Results

In the second experiment, our manipulations of high, and low privacy protection levels again elicited the hypothesized effect. Participants in the low protection condition reported significantly higher beliefs that their

² Of the participants who answered the exit questions, 76.99% indicated they had participated in more than one study and 90.21% reported differences existed between both studies. Results do not differ when we exclude these participants.

responses would be linked back to them ($M_{Low} = .82$, $M_{High} = .42$, $t(370.84) = 8.9794$, $p < .001$, $d = 1.86$) relative to participants in the high-protection condition.

We first evaluated the disclosure rates of participants in the first survey. We found that participants were not more likely to disclose information ($\beta_{High} = .01$, $p = .74$) when they were provided with a high level of protection in the first survey (see Table 4, column 1). Our results were consistent ($\beta_{High} = .01$, $p = .73$) when we included controls for participant demographics (see Table 4, column 2). However, we did not control for the questions' intrusiveness or varying survey designs.

Next, we evaluated disclosure behavior in the second survey of the experiment, in which participants were presented with increasing, decreasing, or identical privacy protection levels compared to the first survey. We first compared participants who had high levels of protection in both surveys with participants who had low levels of protection in both surveys (see Table 4, columns 3 and 4). We included an additional control variable to account for the possibility that high disclosure in the first survey influenced second survey disclosures, using Survey1Sharing, which ranged from a value of 0 (for participants who did not admit to any of the behaviors in Survey 1) to a value of 6 (for participants who admitted to all behaviors in Survey 1). In line with our results for the first survey, we found no effect of high protection versus low protection on disclosure ($\beta_{High} = .01$, $p = .70$) in the second survey (see Table 4, column 3). This result was robust ($\beta_{High} = .01$, $p = .71$) when including controls for participant demographics (see Table 4, column 4). All in all, our results did not provide evidence that changes in objective privacy protection levels influenced actual information disclosure (H1 is not supported). However, the control variable capturing Survey1Sharing turned out to be significant. This pointed toward a person-specific level of disclosure. We will discuss this in Section 6.

Second, we evaluated the impact of relative changes of privacy protection levels on disclosure compared to conditions, in which participants did not perceive an increase or decrease (participants received objectively equivalent privacy protection notices). We found *no* increase in the propensity to disclose information ($\beta_{Increasing} = .02$, $p = .57$) for participants who perceived an increase in protection relative to those whose protections stayed constant. This result was robust when controls for participant demographics were included (see Table 4, columns 5–6). We also found no significant decrease in the overall propensity to disclose ($\beta_{Decreasing} = -.03$, $p = .30$) for participants who perceived a decrease in protection relative to those whose protections stayed constant (see Table 4, column 7). Again, this result was robust when controls for participant demographics were included (see Table 4, column 8). These results suggest that participants' relative perceptions of privacy protection did not impact actual disclosure behavior (H2 is not supported).

4.3 Discussion

Experiment 2 further differentiated the findings from Experiment 1 by investigating actual disclosure settings. However, in contrast to the proposed hypothesis, we did not find any significant impact of either normative or behavioral factors on actual information disclosure (H1 and H2 are not supported). The combined results of Experiment 1 and 2, therefore, indicate that normative factors are stronger in hypothetical disclosure settings than in actual disclosure settings (H3 is supported). The combined results, however, do not demonstrate that behavioral factors are weaker in hypothetical disclosure settings than in actual disclosure settings (H4 is not supported). This phenomenon may be explained by the significant effect of the Survey1Sharing variable, which indicates a person-specific level of disclosure.

Table 4. Experiment 2, Regression Results

Variables	Probability of Admission																
	Our results								Original results								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	
High Protection	0.009 (0.027)	0.009 (0.027)	0.014 (0.036)	0.013 (0.035)					0.0499* (0.0240)	0.0423+ (0.0231)	-0.00336 (0.0278)	0.0001 (0.0278)					
Increasing					0.018 (0.032)	0.02 (0.032)							0.0605* (0.0292)	0.0604* (0.0292)			
Decreasing							-0.034 (0.033)	-0.031 (0.033)							-0.075** (0.0269)	-0.071** (0.0271)	
Intrusive		-		-		-		-		0.076** (0.0178)				-0.111** (0.0259)		-0.086** (0.0288)	
Age		0.003+ (0.002)		0.000 (0.002)		-0.001 (0.002)		-0.003+ (0.002)					0.00139 (0.0016)			0.00057 (0.0016)	
Male		-0.002 (0.027)		-0.084* (0.035)		-0.105* (0.032)		-0.018 (0.033)		0.0493* (0.0232)		0.0512+ (0.0301)		0.0633* (0.0303)		0.0483 (0.0307)	
Survey Design		-		-		-		-		0.0379 (0.0231)				0.00729 (0.0305)		-0.0234 (0.0276)	
Survey 1 Sharing			0.074** (0.011)	0.077** (0.011)	0.081** (0.011)	0.085** (0.011)	0.087** (0.011)	0.086** (0.011)					0.105** (0.0093)	0.095** (0.0102)	0.097** (0.0105)	0.110** (0.0099)	0.109** (0.0103)
Constant	0.858* (0.043)	0.938 (0.059)	0.693** (0.059)	0.720** (0.079)	0.668** (0.06)	0.724** (0.077)	0.711** (0.055)	0.800** (0.071)	0.444** (0.0176)	0.525** (0.0438)	0.0149 (0.0273)	0.0206 (0.0693)	0.0408 (0.0302)	-0.0273 (0.0654)	0.00092 (0.0278)	0.0265 (0.0700)	
Observations	2,063	2,063	981	981	1,948	1,948	2,108	2,108	2,490	2,454	1,164	1,140	1,158	1,140	1,050	1,032	
Number of id	412	412	204	204	196	196	216	216	415	409	194	190	193	190	175	172	

Robust standard errors in parentheses; ** p < 0.01, * p < 0.05, +p < 0.1.; missing data –

5 Experiment 3: Hypothetical and actual information disclosure

5.1 Methodology

Experiment 3 investigated actual and hypothetical disclosure settings simultaneously to confirm that both behavioral and normative factors influence privacy decision making. Again, participants had to participate in two separate surveys, each with a different look and feel. Participants were randomly assigned to one of the eight groups.

The first survey served to set either a high or a low level of privacy protection (as in Experiment 2) but did not include self-disclosure measures. As in the original study, we did not request that participants indicate their disclosure behavior during the first survey, because we did not want actual disclosure to influence disclosures in the second survey. Privacy protection levels were graphically displayed as in Experiment 1 (see Appendix B, Figure B1). Participants had then to rate the level of privacy protection offered in the survey. In the low-protection condition, participants were asked to provide their email address to receive a confirmation code via email for their participation and to increase the perception that answer could be linked to their identity. Identical to the original study, participants had to complete a filler task that separated the first and the second survey. The filler task comprised a 5-minute video about business models and answering questions about the content. A non-privacy filler task enabled participants to encounter two different privacy settings with an extensive delay between both, which better represents real-world privacy scenarios. As in the previous experiments, we included an attention check.

In the second survey, participants were manipulated in such a way that they perceived either an increase, a decrease, or no change in the level of privacy protection compared to the first survey. Participants had then to rate the level of privacy protection offered in the survey. In the low-protection condition, participants were asked to provide their email address to receive a confirmation code via email for their participation. Next, participants were assigned either to the hypothetical or actual disclosure setting. In the actual disclosure setting, participants had to answer five personal and sensitive questions (Acquisti et al., 2012) (see Appendix D, Table D1). In the hypothetical disclosure setting, the questions remained the same; however, participants were asked to imagine participating in a study with certain privacy protection levels provided to the answers. Participants answered a set of questions referring to (un)ethical behaviors and indicated their likelihood of admitting such behaviors. At the end, participants indicated their gender and age. Figure 4 illustrates the process of Experiment 3 in a flow chart.

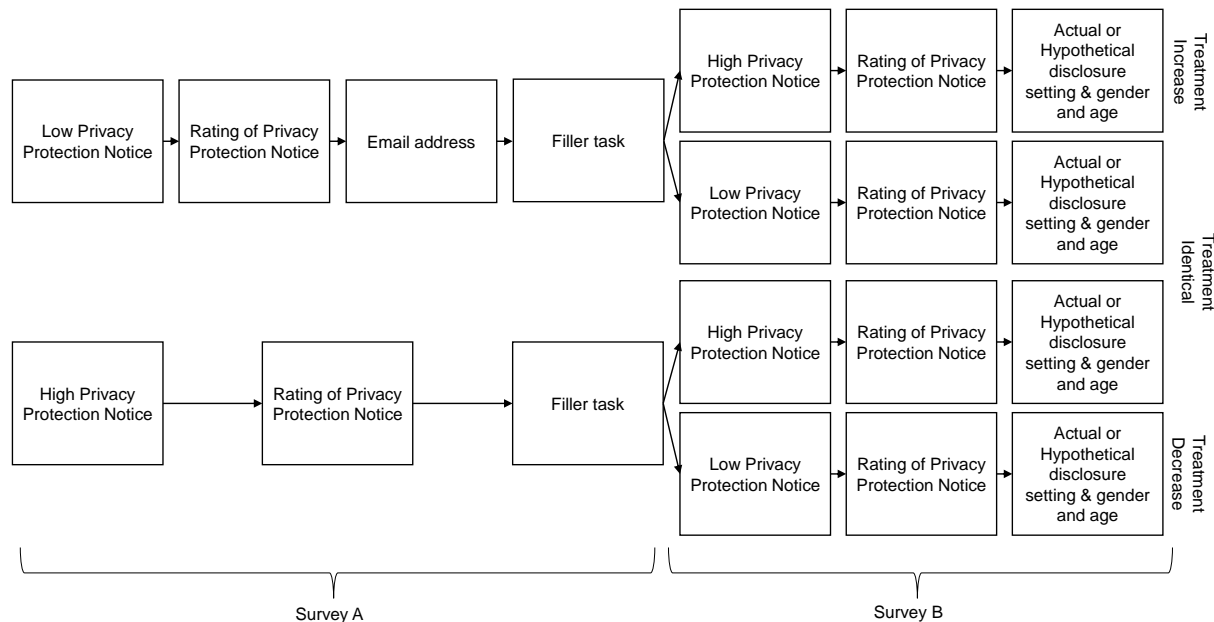


Figure 4: Flow Chart of Experiment 3

Table A3 in Appendix A shows the demographic data of all participants who successfully completed the experiment. In the original study the total sample size for the third experiment was 739 (51.7% were males and a mean age of 29.67, SD of age is 10.1) (Adjerid et al., 2018). Our sample size is 672 (52.38% were males and a mean age of 26.63, SD of age is 10.58).

5.2 Results

In the first study, participants in the high-protection condition rated the study as offering a higher level of privacy protection ($M = 3.76$ vs 2.48 , $SD = 1.27$, $t(672.14) = -15.34$, $p < .001$). We found no consistent results for the ratings of privacy protections in the second study ($M = 3.13$ vs 3.11 , $SD = 1.28$, $t(683.32) = -0.2234$, $p < .82$). We, therefore, could not conclude that our manipulation in the second survey worked as expected. However, the pattern and significance of the results remained the same when excluding manipulation failures, and we therefore report the results of the full sample. We discuss this circumstance in Section 6.

Now, we present the effects on actual and hypothetical disclosure. We first examine participants in the hypothetical settings, where we consider participants to have admitted to the behavior if they responded with either “strongly agree” or “agree” to the question as to whether they would admit to a particular behavior. We do not find statistically significant differences in hypothetical admission rates between those with objectively different (high vs. low) levels of protection (67% versus 52%, $t(622) = -1.197$, $p = .23$). Furthermore, we do not find any significant differences in hypothetical admissions when privacy protections are held objectively constant but relatively decrease (63% versus 57%, $t(461) = 1.44$, $p = .15$) or relatively increase (68% vs. 71%, $t(525) = -0.84$, $p = .40$). We verified that these results are robust to (1) alternative measurements for hypothetical admission, including a continuous measure (i.e., 1–5 on the Likert scale) and (2) considering participants that reported to be uncertain (neither agree nor disagree) as also admitting the behavior.

We confirm these results in a random-effects regression (see Table 5). We find that neither objective differences (high protection) nor relative changes in privacy protection have a significant effect in the hypothetical disclosure setting (see columns 1, 2, and 3). The pattern and significance of the results remain similar when we include participants who failed the manipulation checks. We therefore report the result of the full sample.

Subsequently, we consider participants in the actual-disclosure condition, where participants were shown the same privacy protections and asked the same questions as their counterparts in the hypothetical-disclosure condition. For these participants, we considered an admission as any response to our questions that indicated that the participant engaged in a particular behavior at least once (the same measurement of admission rates was used in the original study). Unlike in the hypothetical context, we find statistically significant differences in actual disclosure behavior between participants with objectively different (high versus low) privacy protections (57% versus 48%, $t(553) = -2.01$, $p = .045$). In line with the hypothetical condition, we find that those who perceived a relative decrease in protection did not disclose significantly less than those who did not perceive a change (45% versus 48%, $t(525) = 0.74$, $p = .46$). Finally, unlike in the hypothetical context, we find that those who perceived a relative increase in protection disclosed significantly more than those who did not perceive a change (57% versus 48%, $t(609) = 2.24$, $p = .03$).

However, the random-effects regression does not support the findings of the t-tests (see Table 6). The regression analysis does reveal a significant effect of objective differences, but the effect becomes insignificant when we exclude participants who failed the manipulation check (see column 1). All other results of Experiment 3 are robust when controlling for failed manipulation checks, and we therefore report the results of the full sample. Furthermore, the regression does not support the initial finding that a relative increase in privacy protection leads to higher levels of actual disclosure (see column 3). However, the regression confirms that a relative decrease in privacy protection does lead to lower levels of actual disclosure (see column 3).

Variables	Probability of Admission					
	Our results			Original results		
	(1)	(2)	(3)	(1)	(2)	(3)
High Protection	0.036 (0.042)			0.0878* (0.0441)		
Decreasing		-0.041 (0.047)			-0.0305 (0.0459)	
Increasing			0.044 (0.044)			0.00185 (0.0433)
Age	-0.004* (0.002)	-0.003 (0.002)	-0.007** (0.002)	-0.00154 (0.00251)	-0.000852 (0.00231)	-0.000954 (0.00276)
Male	0.078+ (0.043)	0.001 (0.046)	0.044 (0.044)	-0.105* (0.0453)	-0.107* (0.0467)	-0.0441 (0.0441)
Constant	0.663** (0.086)	0.754** (0.09)	0.775** (0.09)	0.737** (0.0915)	0.720** (0.0910)	0.718** (0.0967)
Observations	821	712	744	950	910	915
Number of id	179	154	162	190	182	183

Robust standard errors in parentheses; **p < 0.01, *p < 0.05, +p < 0.1.

Variables	Probability of Admission					
	Our results			Original results		
	(1)	(2)	(3)	(1)	(2)	(3)
High Protection	0.097* / 0.05 ¹ (0.044)			0.0552 (0.0410)		
Decreasing		-0.007 (0.047)			-0.108* (0.0476)	
Increasing			-0.069+ (0.042)			-0.0126 (0.0354)
Age	-0.005** (0.002)	-0.005+ (0.003)	0.006** (0.002)	0.00143 (0.00277)	0.00248 (0.00232)	-1.24e-05 (0.00192)
Male	0.137** (0.043)	0.075 (0.048)	0.130** (0.042)	-0.0296 (0.0408)	-0.0826+ (0.0480)	-0.0321 (0.0366)
Constant	0.464** (0.085)	0.562** (0.102)	0.739** (0.093)	0.594** (0.0959)	0.646** (0.0987)	0.695** (0.0694)
Observations	778	696	833	895	810	1,010
Number of id	171	153	180	179	162	202

Robust standard errors in parentheses; **p < 0.01, *p < 0.05, +p < 0.1; ¹ excluding manipulation fails.

5.3 Discussion

In Experiment 3, we examined the simultaneous effects of normative and behavioral factors on actual and hypothetical information disclosure. We found no support for an effect of changes in privacy protection levels on disclosure. Neither objective differences nor relative changes influenced participants' hypothetical disclosure. This result also held true for actual disclosure (H1 and H2 are not supported). Moreover, when comparing the coefficients between hypothetical and actual disclosure, we found no support for the hypothesis that normative factors have a stronger influence on hypothetical intentions compared to actual disclosure (H3 is not supported). Nor did we find evidence that behavioral factors have a weaker influence on hypothetical intentions compared to actual disclosure (H4 is not supported).

Briefly concluding all three experiments, we observe some contradictions between Experiments 1 and 2 and Experiment 3. While Experiment 1 supports H1 and H2, Experiment 3 provides no support for these hypotheses in the hypothetical disclosure setting. Similarly, the comparison of Experiments 1 and 2 supports H3, which is not supported in Experiment 3. However, all experiments demonstrate that H4 is not supported and that neither H1 nor H2 is supported in the actual disclosure setting. Hence, we found mixed support for the argument that normative and behavioral factors can influence hypothetical disclosure and large support for the argument that normative and behavioral factors have no influence on actual disclosure.

6 General Discussion and Conclusion

Our study was carried out as a methodological replication of Adjerid et al. (2018). The original study aimed to fill the void in the IS literature about the simultaneous effect of normative factors (objective differences) and behavioral factors (relative changes) on hypothetical and actual information disclosure. Both studies drew upon literature on consumer privacy decision making and behavioral economics literature regarding reference dependency. However, in contrast to the original study, which showed that relative changes were more pronounced in actual disclosure settings and objective differences were more pronounced in hypothetical disclosure settings, we presented some evidence that normative and behavioral factors influenced hypothetical but not actual disclosure. A comparison of the results is illustrated in Table 7.

		Experiment 1	Experiment 2	Experiment 3	
		Hypothetical choice	Actual choice	Hypothetical choice	Actual choice
H1: Objective Privacy Protection	Orig. study	Support	Mixed support	Support	No support
	Our results	Support	No support	No support	No support
H2: Relative Privacy Protection	Orig. study	No support	Support	No support	Support
	Our results	support	No support	No support	No support
H3: Impact of normative factors	Orig. study	Support		Support	
	Our results	Support		No support	
H4: Impact of behavioral factors	Orig. study	Support		Support	
	Our results	No support		No support	

Our findings point to two areas of discussion. First, the comparison of our results and second, the difference between our results and the results of the original study.

6.1 Comparison of our results

Contrary results between Experiment 1 and Experiment 3

While Experiment 1 investigated the influence of changes in objective and relative levels of privacy protection on hypothetical disclosure, Experiment 3 investigated the influences of objective and relative changes on both actual and hypothetical disclosure. The results of Experiments 1 and Experiment 3 contradict in that Experiment 1 supports H1 (changes in objective levels of privacy protection will affect disclosure) and H2 (one's relative perception of the level of privacy protection will influence individual privacy decision making), whereas Experiment 3 provides no support for these hypotheses in the hypothetical disclosure setting. Therefore, we review the design differences between both experiments to explain the contradiction, in particular as the samples of both experiments do not differ. We argue that five design differences exist. First, Experiment 1 used the same survey design for both surveys, while Experiment 3 used two different survey designs. However, since the survey design had no effect in the original study, we neglect this as potential reason for differentiation. Second, participants in Experiment 1 were manipulated from high or low to medium levels of privacy protection. In contrast, in Experiment 3, participants were manipulated from high or low to high or low levels of privacy protection. Third, Experiment 1 measured objective changes based on disclosure behavior in the first survey, whereas Experiment 3 used disclosure behavior from the second survey. Fourth, Experiment 3 did not request that participants indicate their disclosure behavior in the first survey (to better reflect real-world privacy scenarios) while Experiment 1 did. Fifth, Experiment 3 included a non-privacy related filler task and Experiment 1 did not.

Regarding the contradiction within H1 (changes in objective levels of privacy protection will affect disclosure), we argue that Experiment 3 registered such a high number of manipulation failure (69%³) that it may have rendered the remaining sample too small to identify significant effects. Hence, the contradiction may rather stem from manipulation failure than from the differences between both experiments especially since the process of the objective manipulation is less effected by the differences between both experiments. In other words, if the manipulation would have better worked in Experiment 3, the experiment might have yielded the same supporting result as Experiment 1.

Concerning the contradiction within H2 (one's relative perception of the level of privacy protection will influence individual privacy decision making), we posit that the filler task and missing disclosure behavior in the first survey in Experiment 3 may have caused the result to become insignificant. While the filler task may have led participants to forget the privacy protection level of the first survey, missing disclosure behavior may have amplified this effect since participants were not incentivized to recall the first protection level. In contrast, Experiment 1 showed the two privacy protection levels in rapid succession and requested participants to indicate their disclosure behavior after the first survey which may have helped participants to better remember the first protection level when answering the second disclosure questions.

Contrary results between the results of the comparison of Experiment 1 & 2 and Experiment 3

Experiment 1 investigated the influence of changes in objective and relative levels of privacy protection on hypothetical disclosure, whereas Experiment 2 investigated the influence of both changes on actual disclosure and Experiment 3 investigated both changes on actual and hypothetical disclosure simultaneously. However, while the results of the comparison of Experiments 1 and 2 support H3 (the impact of objective changes will be stronger on hypothetical compared to actual disclosure), the results of Experiment 3 do not. We propose that the rejection of H3 in Experiment 3 is a corollary to the rejection of H1 (objective changes influence disclosure) in Experiment 3. Since H3 is dependent upon the outcome of

³ We found no evidence that participants who failed to understand the manipulation were significantly different from participants passing the manipulation check: difference in means for age of 27.2 (for those who failed) vs. 26.3 (for those who didn't fail), $t(432) = -0.709$, $p = .47$ and 44% female (fails) vs. 50% female (not failed), $X^2(1, 686) = 0.96$, $p = .33$.

H1 and H1 had been rejected due to a high number of manipulation failures, H3 had also been rendered insignificant. We, therefore, argue that the contradiction may stem from manipulation failure rather than from the differences between the three experiments. In other words, if the manipulation had worked better in Experiment 3, the experiment might have yielded the same supporting results as the comparison of Experiments 1 and 2.

Comparison of students and their elderly referrals

The use of students as a sample is potentially problematic. Findings derived from a student sample might not be generalizable to the whole population. To ameliorate these concerns, we used snowball sampling, where the initial group of students was asked to recruit additional people to participate in the study, outside of the student population, to obtain a broader and more representative sample. As we collected as little data as possible to ensure high levels of participants' privacy, we were not able to identify the status of participants and to directly control whether students answer significantly different than referrals. However, we approximated the status of "student" by separating our sample by age in two groups (two times, in groups of older than vs. younger than or equal to 23, 25, and 27 years, respectively). In summary, the t-tests revealed statistically significant but no considerable differences in disclosure for eight out of 15 tests (see Appendix E). While these results provide additional credibility for the use of student samples, we conclude that our results are not entirely generalizable.

6.2 Comparison of our results and the results of the original study

We observed that behavioral factors (relative changes in privacy protection) influence hypothetical rather than actual disclosure. Hence, H4 (behavioral factors will be weaker on hypothetical compared to actual disclosure) is not supported, which stands in direct contrast to the original study supporting H4. Although our results are only valid in Experiments 1 (hypothetical disclosure) and 2 (actual disclosure), we argue that the high amount of manipulation failures in Experiment 3 (hypothetical and actual disclosure) rendered the effect insignificant and that we might have found consistent results among the experiments if the manipulation had succeeded.

We propose two lines of reasoning for the observed difference in our result and that of the original. First, behavioral factors may influence hypothetical disclosure, because (1) the hypothetical context triggers positive attitudes toward disclosing/concealing information (Ajzen et al., 2004) and (2) these positive attitudes foster hypothetical disclosure/concealment. Moreover, participants may fear fewer or even no consequences of their behavior in the hypothetical context and may, therefore, be willing to disclose/conceal more information. Hence, positive attitudes and a lack of consequences may explain why the results indicate that behavioral factors influence hypothetical rather than actual disclosure.

Our second line of reasoning argues that behavioral factors do not influence actual disclosure because participants may have recently developed such a strong disposition to privacy that they are not willing to reveal their actual disclosure behavior no matter the manipulation. By disposition to privacy we refer to "a person's general desire or need for privacy across contexts" (Y. Li, 2014). Such a disposition may have recently emerged and may be more pronounced in our German sample. While Adjerid et al. (2018) collected their data around 2012 and 2016⁴ during which privacy scandals and data protection were less pronounced and discussed in public (Clement, 2019a, 2019b), our samples were exposed to a continuously increasing stream of major privacy breaches and fake news revelations over the last years (e.g. Facebook's influence in the US presidential election 2016 (Cadwalladr & Graham-Harrison, 2018) and in the Brexit referendum (Cadwalladr, 2017)) as well as to public and scholarly debates about GDPR. Hence, our participants may have been more restrictive or even reluctant to disclose actual behavior independent of the level of privacy

⁴ According to the original study the "early analysis of Experiment 2 was published as part of the ACM proceedings from the 2013 Symposium on Usable Privacy and Security" and Experiment 3 has been based on data "of September 2016" (Adjerid et al., 2018).

protection offered due to recent awareness of privacy violations and mistrust toward entities collecting data. Prior work already indicates that increased awareness of privacy violations reduces trust and that trust reduction lowers disclosure (Dinev & Hart, 2006; Malhotra et al., 2004).

In addition to participants' increased disposition to privacy, we argue that the questions may have been too intrusive to elicit true responses about actual behavior. The original study found already a significant effect of intrusiveness on disclosure (Adjerid et al., 2018) as well as prior work (Malhotra et al., 2004) and in combination with high disposition to privacy participants may have decided to conceal their true actual behavior regardless of the manipulation.

Moreover, the original study primarily relied upon American participants and participants from Amazon Mechanical Turk. Both characteristics indicate that the sample of the original study was less privacy sensitive and more prone to information disclosure than our German sample. Not only does the language Americans and Germans use to discuss privacy reflect different ways of conceiving privacy ("privacy" versus "data protection") (Fromholz, 2000), but the divergent levels of uncertainty avoidance in both societies (Hofstede-Insights, 2019; Hofstede, 2001) also indicate that Germans may be more privacy-sensitive. For example, privacy protections in the USA (low on uncertainty avoidance) are mainly based on industry self-regulation, whereas Germany (high on uncertainty avoidance) has substantial laws in place to protect privacy (Bellman, Johnson, Kobrin, & Lohse, 2004). Related research further shows that Germans are more likely than Americans to believe that information provided on Facebook has a higher likelihood of negative outcomes and assume higher damages should these negative outcomes occur (Krasnova & Veltri, 2010).

Returning to Amazon Mechanical Turk, we argue that these participants may be less privacy-sensitive and more prone to disclose (even unethical) information since they disclose this information in return for monetary rewards and may, therefore, feel morally obliged to disclose true behavior (to achieve a good rating) and thus have fewer inhibitions to disclosing actual behavior.

Our first argument (strong individual disposition to privacy) is supported by the significant effect of the survey1sharing variable (approximating disposition to privacy) in Experiment 2 as it indicates that person-specific tendencies toward privacy account for changes in disclosure rather than objective or relative changes in protection. Our second argument (cultural differences regarding privacy) is supported by the fact that our results generally reflect very high privacy levels (low levels of disclosure), compared to those of the original study. In Experiment 1 for example, the original study demonstrates constantly higher baseline disclosure compared to our study (3.63 vs. 3.23, 4.17 vs. 3.67, 3.33 vs. 2.78, 3.77 vs. 3.40). We conclude that our participants have not been willing to indicate their true choice in the actual disclosure setting and thereby rendered normative and behavioral factors insignificant. In other words, at minimum behavioral factors have an impact in the actual setting (Adjerid et al., 2018), but may not be elicitable from privacy-sensitive individuals. This has important implications for scholars relying upon participants to truly report their actual behavior. Our conclusion suggests that these self-reports becomes more and more difficult for privacy-sensitive individuals.

Finally, we partially confirm the findings of Adjerid et al. (2018) that the impact of normative factors (objective changes in privacy protection) is more pronounced in the hypothetical than in the actual disclosure setting. That is, the results of Experiment 1 (hypothetical disclosure) and 2 (actual disclosure) support H3 (normative factors will be stronger on hypothetical compared to actual disclosure) but the results of Experiment 3 (hypothetical and actual disclosure) do not support H3. However, since Experiment 3 suffers from high manipulation fails, we need to interpret the results carefully and we therefore propose that consistent results may have merged from the experiments if the manipulation had succeeded.

7 Limitations and Future Research

Our study possesses several limitations. First, as Adjerid et al. (2018) pointed out, this work investigated specific factors within the normative and behavioral perspective. However, other biases such as framing effects, isolation effects, or bandwagon effects may lead to different findings. We, therefore, encourage future work to extend our manipulations by exploring how other cognitive biases affect privacy decision making. Second, although we tried to reach out to different sociodemographic groups, our study mainly

comprises students. Our work is therefore not representative of the entire German population. We suggest future work to engage in more representative studies to assess the extent, to which different sociodemographic groups are prone to manipulation and whether some groups may need more regulatory protection than others (in case biases are used to harm consumers, e.g. through less protective default settings). Third, our results suggest that an individual's disposition to privacy and their cultural background inhibited manipulation in the actual disclosure setting, but we did not directly control for those aspects. Hence, it seems fruitful to explore the moderating effect of disposition to privacy and uncertainty avoidance (as a more specific subdimension of culture) in future studies. Finally, while our manipulations succeeded in Experiment 1 and 2, the manipulations did not work well in Experiment 3 and therefore need to be interpreted carefully.

Acknowledgments

We sincerely would like to thank the editors and reviewers of the Association for Information Systems Transactions on Replication Research for their valuable contribution to improving the quality of this paper.

References

- Acquisti, A., John, L. K., & Loewenstein, G. (2012). The impact of relative standards on the propensity to disclose. *Journal of Marketing Research*, 49(2), 160-174.
- Acquisti, A., John, L. K., & Loewenstein, G. (2013). What is privacy worth? *The Journal of Legal Studies*, 42(2), 249-274.
- Adjerid, I., Peer, E., & Acquisti, A. (2018). Beyond the privacy paradox: Objective versus relative risk in privacy decision making. *MIS Quarterly*, 42(2), 465-488.
- Ajzen, I., Brown, T. C., & Carvajal, F. (2004). Explaining the discrepancy between intentions and actions: The case of hypothetical bias in contingent valuation. *Personality and social psychology bulletin*, 30(9), 1108-1121.
- Ansari, A., & Mela, C. F. (2003). E-Customization. *Journal of Marketing Research*, 40(2), 131-145.
- Bellman, S., Johnson, E. J., Kobrin, S. J., & Lohse, G. L. (2004). International differences in information privacy concerns: A global survey of consumers. *The Information Society*, 20(5), 313-324.
- Cadwalladr, C. (2017). The great British Brexit robbery: How our democracy was hijacked. *The Guardian*. Retrieved from <https://www.theguardian.com/technology/2017/may/07/the-great-british-brexit-robbery-hijacked-democracy>.
- Cadwalladr, C., & Graham-Harrison, E. (2018). Revealed: 50 million Facebook profiles harvested for Cambridge Analytica in major data breach. *The Guardian*. Retrieved from <https://www.theguardian.com/news/2018/mar/17/cambridge-analytica-facebook-influence-us-election>.
- Clement, J. (2019a). Annual number of data breaches and exposed records in the United States from 2005 to 2018 *Statista*. Retrieved from <https://www.statista.com/statistics/273550/data-breaches-recorded-in-the-united-states-by-number-of-breaches-and-records-exposed/>.
- Clement, J. (2019b). Number of compromised data records in selected data breaches as of May 2019. *Statista*. Retrieved from <https://www.statista.com/statistics/290525/cyber-crime-biggest-online-data-breaches-worldwide/>.
- Dennis, A., & Valacich, J. (2015). A Replication Manifesto. *AIS Transactions on Replication Research*, 1, 1-5.
- Dinev, T., & Hart, P. (2006). An extended privacy calculus model for e-commerce transactions. *Information Systems Research*, 17(1), 61-80.

- Farahat, A., & Bailey, M. C. (2012). *How effective is targeted advertising?* Paper presented at the Proceedings of the 21st International Conference on World Wide Web.
- FeldmanHall, O., Mobbs, D., Evans, D., Hiscox, L., Navrady, L., & Dalgleish, T. (2012). What we say and what we do: The relationship between real and hypothetical moral choices. *Cognition*, *123*(3), 434-441.
- Fromholz, J. M. (2000). The European Union data privacy directive. *Berkeley Technology Law Journal*, *15*(1), 461-484.
- Herrmann, P. N., Kundisch, D. O., & Rahman, M. S. (2014). Beating irrationality: Does delegating to IT alleviate the sunk cost effect? *Management Science*, *61*(4), 831-850.
- Hofstede-Insights. (2019). Country comparison. Retrieved from <https://www.hofstede-insights.com/country-comparison/germany,the-usa/>.
- Hofstede, G. (2001). *Culture's Consequences: Comparing Values, Behaviors, Institutions and Organizations across Nations* (Vol. 2). Thousand Oaks, California: Sage Publications.
- John, L., Acquisti, A., & Loewenstein, G. (2011). Strangers on a plane: Context-dependent willingness to divulge sensitive information. *Journal of Consumer Research*, *37*(5), 858-873.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*(2), 263-291.
- Kang, M. J., & Camerer, C. F. (2013). fMRI evidence of a hot-cold empathy gap in hypothetical and real aversive choices. *Frontiers in Neuroscience*, *7*(Article 104), 1-16.
- Keith, M. J., Babb, J. S., & Lowry, P. B. (2014). A longitudinal study of information privacy on mobile devices. Paper presented at the *47th Hawaii International Conference on System Sciences*, Waikoloa, Hawaii.
- Knetsch, J. L., Tang, F.-F., & Thaler, R. H. (2001). The endowment effect and repeated market trials: Is the Vickrey auction demand revealing? *Experimental Economics*, *4*(3), 257-269.
- Krasnova, H., & Veltri, N. F. (2010). Privacy calculus on social networking sites: Explorative evidence from Germany and USA. Paper presented at the *43rd Hawaii International Conference on System Sciences*.
- LaPiere, R. T. (1934). Attitudes vs. actions. *Social Forces*, *13*(2), 230-237.
- Li, H., Sarathy, R., & Xu, H. (2011). The role of affect and cognition on online consumers' decision to disclose personal information to unfamiliar online vendors. *Decision Support Systems*, *51*(3), 434-445.
- Li, Y. (2014). The impact of disposition to privacy, website reputation and website familiarity on information privacy concerns. *Decision Support Systems*, *57*, 343-354.
- Loewenstein, G. (2000). Emotions in economic theory and economic behavior. *American Economic Review*, *90*(2), 426-432.
- Malhotra, N. K., Kim, S. S., & Agarwal, J. (2004). Internet users' information privacy concerns (IUIPC): The construct, the scale, and a causal model. *Information Systems Research*, *15*(4), 336-355.
- Mullainathan, S., & Thaler, R. H. (2000). Behavioral economics. *National Bureau of Economic Research Working Paper Series*, No. 7948.
- Neumann, J. v., & Morgenstern, O. (1944). *Theory of Games and Economic Behavior*. Princeton, New Jersey: Princeton University Press.
- Viswanathan, S., Kuruzovich, J., Gosain, S., & Agarwal, R. (2007). Online infomediaries and price discrimination: Evidence from the automotive retailing sector. *Journal of Marketing*, *71*(3), 89-107.

Appendix A: Sample statistics

Table A1. Experiment 1, Demographic data		
	Responses (n)	
Total responses	331	
Did not finish	96	
Responses	235	
Failed attention check	60	
Group	Low privacy protection	High privacy protection
	Sex	
Male	50	38
Female	43	44
Total	93	82
	Age (Quantile)	
0%	19	18
25%	21	21
50%	23	23
75%	26	25
100%	55	64

Table A2. Experiment 2, Demographic data				
	Responses (n)			
Total responses	541			
Did not finish	129			
Responses	412			
Failed attention check	120			
Group	High, High	High, Low	Low, High	Low, Low
	Sex			
Male	44	58	49	54
Female	48	46	55	58
Total	92	104	104	112
	Age (Quantile)			
0%	18	18	18	18
25%	21	20	20	20
50%	23	23	23	23
75%	26	25	26	25
100%	51	66	67	91

Table A3. Experiment 3, Demographic data								
	Responses (n)							
Total responses	949							
Did not finish	277							
Responses	672							
Failed attention check	172							
Group (l=Low, h=High, a=actual, b=hypothetical)	h, h, a	h, h, b	h, l, a	h, l, b	l, h, a	l, h, b	l, l, a	l, l, b
	Sex							
Male	56	41	42	35	47	44	39	48
Female	39	45	40	27	44	33	44	48
Total	95	86	82	62	91	77	83	96
	Age (Quantile)							
0%	18	18	18	18	18	18	18	18
25%	21	21	22	22	20	20	20	20
50%	24	24	24	24	23	23	23	23
75%	28	26	27	27	26	25	26	26
100%	74	60	57	83	65	62	62	75

Appendix B: Experiment 1

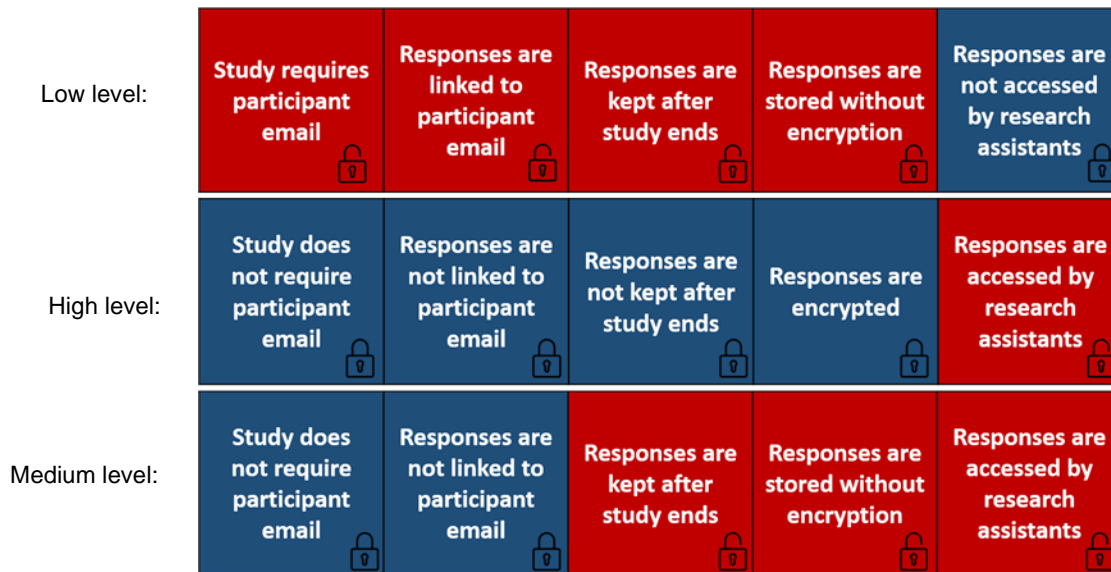


Figure B1. Low, High, Medium privacy level

The locks should indicate the privacy focus, an open lock means lower privacy and closed lock means higher privacy level.

Table B1. Manipulation checks and protection recall		
Measure	Description	Response scale
Privacy Concern	I would be concerned about my privacy if I was participant in this upcoming survey A/B.	Strongly Agree – Strongly Disagree [5 scale]
Protection Satisfaction	I am satisfied with the protections provided in this upcoming survey A/B.	Strongly Agree – Strongly Disagree [5 scale]
Harm Perception	I would be concerned that my responses in this upcoming survey A/B could be used to harm me.	Strongly Agree – Strongly Disagree [5 scale]
Protection Recall 1	Does survey A/B require a valid email address?	Yes, No
Protection Recall 2	The responses in survey A/B are linked to my email.	Yes, No
Protection Recall 3	My responses are kept after the end of survey A/B.	Yes, No
Protection Recall 4	My responses are encrypted in survey A/B.	Yes, No
Protection Recall 5	My responses in survey A/B will be accessed by a research assistant.	Yes, No

As in the original study the following text was given to the participants: Imagine you are taking study A/B. How likely are you to truthfully answer the following questions?

Table B2. Hypothetical Questions	
Description	Response scale [5 scale]
What is your annual income?	Very Unlikely – Very Likely
What is your sexual orientation?	Very Unlikely – Very Likely
What is your address?	Very Unlikely – Very Likely
What is your phone number?	Very Unlikely – Very Likely
What is your view on gay rights?	Very Unlikely – Very Likely
Have you every downloaded a pirated song?	Very Unlikely – Very Likely
Have you ever flirted with someone other than your partner or spouse?	Very Unlikely – Very Likely
Have you ever used drugs of any kind (e.g., weed, heroin, crack)?	Very Unlikely – Very Likely
Have you ever looked at pornographic material?	Very Unlikely – Very Likely
Have you ever made up a serious excuse, such as a grave illness or death in the family, to get out of doing something?	Very Unlikely – Very Likely

Appendix C: Experiment 2

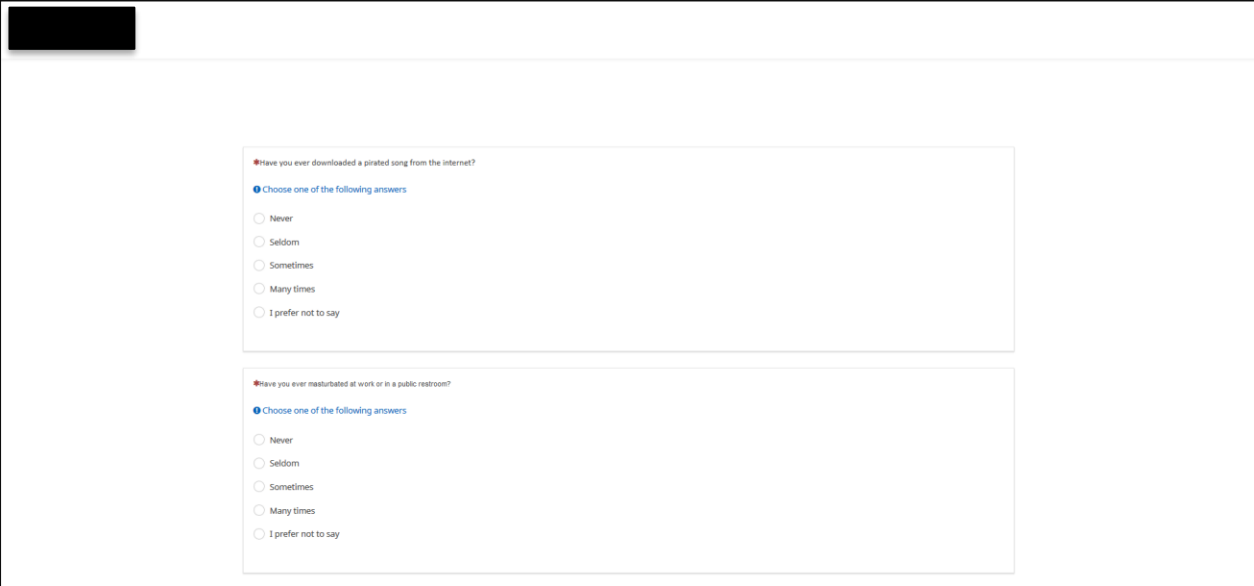
Privacy notice	Text
High	The analysis for this study requires that your responses are stored using a randomly assigned ID. All other information that could potentially be used to identify you (email, zip code, etc.) will be stored separately from your responses. As such, your responses to the following set of questions cannot be directly linked back to you.
Low	The analysis for this study requires that your responses are stored using your email. As such, your responses to the following set of questions may be directly linked back to you.

Measure	Description	Response scale
Protection Recall 1	Does survey A/B requires a valid email address?	Yes, No
Protection Recall 2	The responses in survey A/B are linked to my email.	Yes, No
Privacy Concern	I am concerned about my privacy in this survey.	Strongly Agree – Strongly Disagree [5 scale]
Protection Satisfaction	I am satisfied with the protections provided in this survey.	Strongly Agree – Strongly Disagree [5 scale]
Harm Perception	I am concerned that my responses in this survey could be used to harm me.	Strongly Agree – Strongly Disagree [5 scale]

The scale for the following questions ranged from never to many times, with an additional option I prefer not to say.

Description	Study
Have you ever downloaded a pirated song from the internet?	A
While in a relationship, have you ever flirted with somebody other than your partner?	A
Have you ever masturbated at work or in a public restroom?	A
Have you ever fantasized about having violent nonconsensual sex with someone?	A
Have you ever tried to gain access to someone else's (e.g., a partner, friend, or colleague's) email account?	A
Have you ever looked at pornographic material?	A
Have you ever used drugs of any kind (e.g., weed, heroin, crack)?	B
Have you ever let a friend drive after you thought he or she had had too much to drink?	B
Have you ever made up a serious excuse, such as grave illness or death in the family, to get out of doing something?	B
Have you ever had sex in a public venue (e.g., restroom of a club, airplane)?	B
Have you ever, while an adult, had sexual desires for a minor?	B
Have you ever had a fantasy of doing something terrible (e.g., torture) to someone?	B

Table C4. Exit questions	
Description	Response scale
The confidentiality protections in this study [were the same as, increased relative to, decreased relative to] the confidentiality protections in the prior study.	[Strongly Agree – Strongly Disagree] [5 scale]
As part of this hit, you participated in:	[One Study, Two Separate Studies, Three Separate Studies]
What are the differences between the first and second study?	[No Difference, Different Questions, Different Confidentiality Protections, Different Purpose]



The screenshot shows a survey interface with two questions, each with five radio button options. The first question is: "Have you ever downloaded a pirated song from the internet?" with options: "Never", "Seldom", "Sometimes", "Many times", and "I prefer not to say". The second question is: "Have you ever masturbated at work or in a public restroom?" with the same five options. The interface includes a "Choose one of the following answers" instruction for each question.

Figure C1. Design of Survey A

The image shows a screenshot of a survey interface with a dark blue background. At the top, there is a black rectangular redaction. Below it, the text "Please answer the following questions." is centered. The survey contains two questions, each in a white rounded rectangle. The first question asks, "Have you ever used drugs of any kind (e.g., weed, heroin, crack)?" and offers five radio button options: "Never", "Seldom", "Sometimes", "Many times", and "I prefer not to say". The second question asks, "Have you ever let a friend drive after you thought he or she had had too much to drink?" and offers the same five radio button options. A small red asterisk is visible in the top left corner of each question box.

Figure C2. Design of Survey B

Appendix D: Experiment 3

Actual response scale: Never - many times, additional option: I prefer not to say
Hypothetical response scale: [Definitely no - Definitely yes], 5 points

Table D1. Actual and hypothetical questions
Description
Have you ever downloaded a pirated song from the internet?
While in a relationship, have you ever flirted with somebody other than your partner?
Have you ever looked at pornographic material?
Have you ever used drugs of any kind (e.g., weed, heroin, crack)?
Have you ever made up a serious excuse, such as grave illness or death in the family, to get out of doing something?

Appendix E: Generalizability of results

As there was no identifier for students vs. their referrals, we approximated the status of participants by splitting them in two groups at various split points by their age. Table E1 reports the disclosure of the two groups including results from t-tests for all three experiments for all of these split points.

	Mean of the group: younger/equal	Mean of the group: older	DF	t	p
Experiment 1 with split at age=23, Survey 1	3.087	3.040	232.708	-0.390	0.697
Experiment 1 with split at age=25, Survey 1	3.137	2.881	133.972	-2.037	0.044
Experiment 1 with split at age=27, Survey 1	3.124	2.772	60.759	-2.397	0.020
Experiment 1 with split at age=23, Survey 2	2.981	2.854	229.349	-0.978	0.329
Experiment 1 with split at age=25, Survey 2	2.976	2.775	120.233	-1.401	0.164
Experiment 1 with split at age=27, Survey 2	2.986	2.590	55.531	-2.293	0.026
Experiment 2 with split at age=23, Survey 1	2.669	2.530	361.574	-2.009	0.045
Experiment 2 with split at age=25, Survey 1	2.634	2.538	152.856	-1.144	0.254
Experiment 2 with split at age=27, Survey 1	2.643	2.423	76.338	-2.087	0.040
Experiment 2 with split at age=23, Survey 2	2.417	2.249	345.792	-2.683	0.008
Experiment 2 with split at age=25, Survey 2	2.351	2.332	148.729	-0.252	0.801
Experiment 2 with split at age=27, Survey 2	2.354	2.301	72.758	-0.519	0.606
Experiment 3 with split at age=23	2.208	2.111	680.658	-1.656	0.098
Experiment 3 with split at age=25	2.218	2.025	378.169	-3.023	0.003
Experiment 3 with split at age=27	2.227	1.910	227.395	-4.515	0.000

About the Authors

Sebastian Hermes is a research associate and Ph.D. student at the Chair for Information Systems at Technical University of Munich (TUM), Germany. He holds a Master's degree in Entrepreneurship from the University of Liechtenstein and a Bachelor's degree from the Baden-Wuerttemberg Cooperative State University. Sebastian has worked five years at Roche Pharma and co-founded Thinkfield. His work has appeared in *Electronic Markets* and in refereed conference proceedings such as the *International Conference on Information Systems (ICIS)* and the *European Conference on Information Systems (ECIS)*.

Luis Hillebrand is a master student in Management & Technology at the Technical University of Munich (TUM), Germany. He holds Bachelor's degrees in Management & Technology from TUM and in Philosophy from Ludwig Maximilian University of Munich (LMU), Germany. He presented his research at the *Second Transatlantic Conference on Data & Ethics* in Vienna and the *Buffalo Experimental Philosophy Conference*.

Jan Bauer is a master student in Information Systems at the Technical University of Munich (TUM), Germany. He holds a Bachelor's degree in Information Systems from the University of Applied Sciences Albstadt-Sigmaringen. Jan has working experience in multiple companies like Amazon Web Services, Lufthansa, and Porsche. His work has appeared in the refereed conference proceeding of the *International Conference on Intelligent Decision Technologies*.

Markus Böhm is research group leader at the Chair for Information Systems at Technical University of Munich (TUM), Germany. He graduated in Business & Information Systems Engineering from Friedrich-Alexander University Erlangen-Nürnberg (FAU), and holds a doctoral degree in Information Systems from TUM. His research focus is on mergers & acquisitions, business model innovation and digital transformation. His work has appeared in *Management Information Systems Quarterly Executive*, *Electronic Markets*, *Communications of the Association of Information Systems* and several refereed conference proceedings, including *ECIS* and *ICIS*.

Helmut Krcmar is a Chair Professor of Information Systems at Technical University of Munich (TUM), Germany. Before 2002, he was Chair for Information Systems, University of Hohenheim, Stuttgart. Helmut is an AIS Fellow and has served the IS community in many roles, including as President of the Association for Information Systems. His research interests include information and knowledge management, service management, business process management, and business information systems. His work has appeared in *Management Information Systems Quarterly*, *Journal of Management Information Systems*, *Journal of Strategic Information Systems*, *Journal of Management Accounting Research*, *Journal of Information Technology*, *Information Systems Journal*, and *Business & Information Systems Engineering*.

Copyright © 2020 by the Association for Information Systems. Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and full citation on the first page. Copyright for components of this work owned by others than the Association for Information Systems must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or fee. Request permission to publish from: AIS Administrative Office, P.O. Box 2712 Atlanta, GA, 30301-2712 Attn: Reprints or via e-mail from ais@aisnet.org.